APPLIED MATHEMATICS TECHNICAL REPORT

BOEING COMPUTER SERVICES ETA-TR-40

Experiences in Solving Large Eigenvalue Problems on the CRAY X-MP

R. G. Grimes
J. G. Lewis
H. D. Simon

November, 1986

Applied Mathematics Unit
Engineering Technology Applications Division
G-6500, M/S 7L-21
P.O. Box 24346
Seattle, Washington 98124-0346

EXPERIENCES IN SOLVING LARGE EIGENVALUE PROBLEMS ON THE CRAY X-MP

R. G. GRIMES
J. G. LEWIS
H. D. SIMON

Boeing Computer Services
Seattle, WA

INTRODUCTION

The Applied Mathematics Unit of Boeing Computer Services has been involved with solving large eigenvalue problems since its inception. In the last seven years parallel advances in both computer architecture and numerical algorithms have changed the class of problems that we now see. The previously hard problems are now trivial. Solutions to the here-to-fore intractible problems are now required. This paper will give an overview of two application areas which generate large eigenvalue problems and the block shifted Lanczos algorithm used in their solution. The computational aspects of the algorithm are outlined. Some performance results with some historical perspective are given.

SOURCE OF EIGENVALUE PROBLEMS

The large eigenvalue problems seen by the authors come from two quite different areas: structural engineering and material science. Structural engineering packages (e.g. MSC/NASTRAN) use a finite element model of the structure to analyze how the structure reacts to forces such as seismic activity and wind loading. Two types of analysis, vibrational and buckling, generate large sparse symmetric generalized eigenproblems.

A vibration analysis requires the computation of a small, relative to the problem order, number of eigenvalues and their corresponding eigenvectors of the problem

$$KX = MX\Lambda,$$

where K is the finite element stiffness matrix, M is the finite element mass matrix, Λ and X are the natural frequencies and normal mode shapes which

are to be computed. The results are then used to approximate the solution of a system of ordinary differential equations

$$M\ddot{x}(t) + Kx(t) = 0.$$

The buckling analysis is the determination of whether a structure buckled under an applied force. Assuming that stiffness depends on the direction in which the force is applied and using a first order Taylor's expansion of the nonlinear stiffness matrix the buckling equation is

$$K\Phi = -\mu K_\delta(f)\Phi,$$

where K is the stiffness matrix, $K_\delta(f)$ is the differential stiffness matrix at the applied force f, Φ is the buckling load and μ is the factor of safety. Here the eigenproblem is to compute the eigenvalue with smallest magnitude. If it is less than 1 the structure will buckle under the applied force. If it is greater than 1 the structure will withstand the applied force with a safety margin given by the eigenvalue. In either case the eigenvector gives the state of the structure under the applied force.

The eigenvalue problems in structural engineering can be quite large with as many as 40000 degrees of freedom. Yet the matrices have only a few nonzeroes per row. The stiffness matrix K is sparse and usually positive semidefinite (it can be indefinite). The nonzero patterns of M and K_δ are usually subsets of K. For the vibration problem M is symmetric positive semidefinite. K_δ, in the buckling problem, is symmetric indefinite. In both problems, the physically meaningful eigenvalues are poorly separated as they tend to be near zero. Yet there are usually some very large, if not infinite, eigenvalues in the spectrum. The vibration problem can have eigenvalues with very high multiplicities.

The eigenproblems arising in material science are quite different. They are dense symmetric generalized eigenproblems of the form

$$AX = BX\Lambda,$$

arising from applying a Rayleigh-Ritz variational model to solve Schrödinger's equation for the quantum mechanical wave function [2]. The eigenproblems are generated in an iterative process where the eigensolution of one step is used to formulate the next eigenproblem. Typically the lowest ten percent of the eigenvalues and eigenvectors are required. The problem order is dynamic with typical sizes near 1000 but it is not unreasonable to consider problems on the order of 2000 and 3000.

LANCZOS ALGORITHM

In recent years the Lanczos algorithm has become the method of choice for large sparse generalized eigenvalue problems. A block shifted Lanczos algorithm has been implemented by the authors and documented in detail in [3]. For a discussion of the basic issues involving the practical use of the Lanczos algorithm for eigenvalue computations see [4].

Here we will only discuss some aspects of the block shifted Lanczos algorithm, as they are relevant for the understanding of the numerical results. The simple, unshifted Lanczos algorithm for the symmetric eigenvalue problem projects the original matrix to a tridiagonal matrix whose eigenvalues approximate the original eigenvalues. As more iterations are taken the approximate eigenvalues become more accurate with the eigenvalues at the extreme ends of the spectrum converging first. The Lanczos algorithm does not modify the original matrix. The only computation involving the matrix is a matrix vector product.

The basic Lanczos algorithm must be modified in several ways to handle the symmetric generalized eigenproblem. One modification is to switch to a block Lanczos algorithm which operates on several vectors simultaneously and reduces the original matrix to block tridiagonal form. Even though convergence becomes slower for larger blocksizes, a block algorithm can become overall more efficient in terms of execution time. Furthermore a block algorithm can easily handle multiple eigenvalues with multiplicity less than or equal to the block size.

Another modification is the use of a a shifted Lanczos algorithm. The shifted algorithm is related to inverse iteration in that the eigenvalues which are close to the chosen shift values converge most rapidly. This allows the Lanczos algorithm to compute approximations of many eigenvalues near the shift. The shifted Lanczos algorithm does require a factorization of $K-\sigma M$ $(K-\sigma K_\delta, A-\sigma B)$ along with the ability to perform forward and backsubstitution operations with the factored matrix.

The implementation of the block shifted Lanczos algorithm [3] also employs an automatic shifting strategy. The user specified computational interval is searched for eigenvalues using judiciously placed shifts. The use of the inertia counts [3,5] assures that no eigenvalues are missed.

The block shifted Lanczos algorithm [3] was developed for sparse problems arising in structural engineering. Its development was funded by the MacNeal-Schwendler Corporation for inclusion in MSC NASTRAN, a structural engineering analysis package. The Lanczos code was developed with

a modular structure to allow it to be used in other applications (see Figure 1). At Boeing Computer Services, Lanczos has been connected with an extended version of SPARSPAK for sparse matrix operations and with LINPACK [1] for dense matrix operations.

In addition to the three matrix operations, factorization, solution, and matrix multiplication, there are several other critical areas that consume significant amounts of compuational time. They are orthogonalization steps within the Lanczos recurrence, the compuation of eigenvectors, Modified Gram- Schmidt, and the eigensolution of the block tridiagonal matrix.

Performance Results

Table 1 lists the performance of Lanczos on the Cray X-MP/24 at Boeing Computer Services on 7 symmetric generalized eigenproblems. These numbers are rather bland except when placed in an historical perspective. Platzman's problems are finite difference discretizations of an oceanographic model for tidal movements. They were formulated in the mid 1970s by G. W. Platzman and proved to very difficult eigenproblems as one of their characteristics is that all of the eigenvalues appear as doubletons. In the 70s the eigenvalues in the interval .0001 to .024 were computed with great difficulty. The eigenvalues in the interval .000025 to .0001 were also of interest but impossible to compute. With the block shifted Lanczos algorithm the solution to these problems took only a few seconds of computer time and required no user intervention for their solution.

The next three problems are vibrational analyses from structural engineering. The Reactor Containment Floor was solved in 1981 using a single vector unshifted Lanczos code on the Cray-1S in 36 seconds. The speed-up seen on the Cray X-MP is attributal to the hardware improvements between the two computers. The Sports Arena also was solved in 1981 with the above mentioned code in 27 seconds. The speed-up seen here is attributal to both the hardware improvements and to the algorithmic improvements of block shifted Lanczos. The Columbia Center problem arose from a full scale finite element model of the 76 story skyscraper in Seattle of the same name. The model was abandoned in 1978 by engineers. It was resurrected in 1981 to test the older Lanczos code. It proved to be infeasible then because of the memory restriction of 2 Million words on our Cray-1S. It was solved for the first time in 1986 using the new algorithm on the X-MP. Unfortunately the finite element model was never completed by the engineers and the eigenvalue problem had 115 zero eigenvalues (i.e. rigid body modes). Block shifted Lanczos demonstrated remarkable robustness in computing all of these 115 eigenvalues so that it would not miss any of the 10 lowest that were requested.

The 767 Bulkhead problem was to compute the first buckling mode of the rear bulkhead. The engineers on the project formulated the finite element model in 1985 and were having difficulty with their statics analysis. The buckling analysis demonstrated a error in their finite element model which enabled them to complete their statics modeling (which does not require an eigensolution) in a timely fashion. Without the buckling analysis they would have had to painstakingly review their finite element model with several thousand elements to find their modeling error.

The material science problem was the largest dense eigenproblem that could be stored using an in-core factorization on the X- MP/24 at BCS. The 47.7 seconds using Lanczos compares to 108.5 seconds using a packed symmetric generalized eigenvalue path in EISPACK3. This path did not exist and the authors created the missing components to allow this mode of solution.

A FLOWTRACE analysis for the distribution of computational time was performed for a wide range of eigenproblems. Depending on the eigenproblem the computational intensive parts of the algorithm were in the sparse matrix reordering; matrix (both sparse and dense) factorization, solution, and multiplication; the tridiagonal matrix analyze step; and the orthogonalization steps. Except for the tridiagonal matrix analyze step, the remainder of the computation were based on computational kernels in VectorPak [5], a package of highly optimized computational kernels for scientific and engineering applications on Cray computers.

An example of the performance improvements due to the Cray X-MP would be the improvement of execution for the sparse matrix factorization of the shifted matrix. The factorization is based on sparse Gaussian elimination which requires an inner loop with indirect addressing. Table 2 presents the execution time of the factorization on two of the matrices discussed earlier. The times have been made relative to the most efficient factorization using a VectorPak kernel which utilizes the hardware scatter/gather feature of the X-MP. The other times are using CFT 1.14 or VectorPak with or without hardware gather/scatter.

SUMMARY

The block shifted Lanczos algorithm coupled with the Cray X-MP provides the most effective solution of large symmetric generalized eigenproblems for both sparse and dense matrices. For sparse matrices the hardware gather/scatter feature of the X- MP make an important reduction in execution time for the matrix operations.

REFERENCES

[1] J.J. Dongarra, C.B. Moler, J.R. Bunch, and G.W. Stewart, *LINPACK User's Guide*, SIAM, Philadelphia, 1979.

[2] R. Grimes, H. Krakauer, J. Lewis, H. Simon, and S. Wei, The Solution of Large Dense Generalized Eigenvalue Problems on the Cray X-MP/24 with SSD, to appear in *J. Comp. Physics* (1987).

[3] R. Grimes, J. Lewis, and H. Simon, *The Implementation of a Block Shifted and Inverted Lanczos Algorithm for Eigenvalue Problems in Structural Engineering*, ETA-TR-39, Boeing Computer Services (1986).

[4] B. Parlett, *The Symmetric Eigenvalue Problem*, Prentice-Hall (1980).

[5] *VectorPak Users Manual*, Boeing Computer Services, Document No. 20460-0501 (1986).

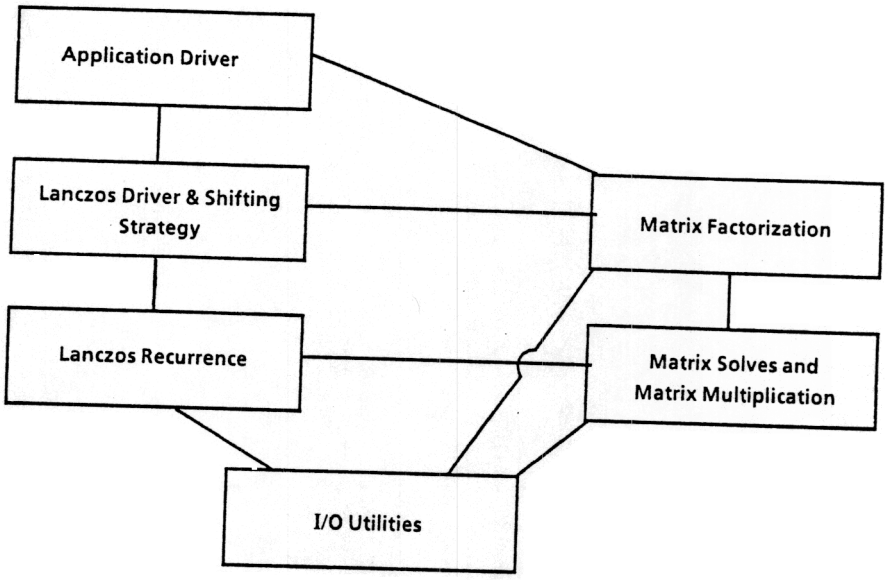

Figure 1. Structure of Lanczos Code

Problem Title	*Order*	*Eigenproblem*	*CP Time (secs)*
Platzman's Small Problem	362	164 e.v. in interval	8.8
Platzman's Large Problem	1919	648 e.v. in interval	93.9
Reactor Containment Floor	1922	200 smallest e.v.	26.8
Sports Arena	3562	10 smallest e.v. (118 computed)	9.5
767 Bulkhead	13992	1 buckling mode (modeling error)	
Material Science	1496	30 e.v. in interval	

Table 1. Performance Results on the Cray X-MP

Implementation	*Sports Arena* (N=3562)	*Columbia Center* (N=15439)
CFT 1.14 w/out g/s	6.3	8.7
VectorPak w/out g/s	2.5	3.1
CFT 1.14 with g/s	1.1	1.2
VectorPak with g/s	1.0	1.0

Table 2. Impact of Hardware Gather/Scatter on Sparse Gaussian Elimination. (CP times normalized so that VectorPak with g/s = 1.0)